arXiv:1402.2073v1 [cs.IR] 10 Feb 2014

Mining Images in Biomedical Publications: Detection and Analysis of Gel Diagrams

Tobias Kuhn,*[1] Mate Levente Nagy,[3] ThaiBinh Luong,[2] and Michael Krauthammer[2,3]

[1]Department of Humanities, Social and Political Sciences, ETH Zurich, Switzerland. [2]Department of Pathology, Yale University School of Medicine, New Haven, CT, USA. [3]Program for Computational Biology and Bioinformatics, Yale University, New Haven, CT, USA.

Email: Tobias Kuhn* – kuhntobias@gmail.com; Mate Levente Nagy — mate.nagy@yale.edu; ThaiBinh Luong – thaibinh@gmail.com; Michael Krauthammer – michael.krauthammer@yale.edu;

*Corresponding author

Abstract

Authors of biomedical publications use gel images to report experimental results such as protein-protein interactions or protein expressions under different conditions. Gel images offer a concise way to communicate such findings, not all of which need to be explicitly discussed in the article text. This fact together with the abundance of gel images and their shared common patterns makes them prime candidates for automated image mining and parsing. We introduce an approach for the detection of gel images, and present a workflow to analyze them. We are able to detect gel segments and panels at high accuracy, and present preliminary results for the identification of gene names in these images. While we cannot provide a complete solution at this point, we present evidence that this kind of image mining is feasible.

Introduction

A recent trend in the area of literature mining is the inclusion of images in the form of figures from biomedical publications [1–3]. This development benefits from the fact that an increasing number of scientific articles are published as open access publications. This means that not just the abstracts but the complete texts including images are available for data analysis. Among other things, this enabled the development of query engines for biomedical images like the Yale Image Finder [4] and the BioText Search Engine [5]. Below, we present our approach to detect and access gel diagrams. This is an extended version of a previous workshop paper [6].

As a preparatory evaluation to decide which image type to focus on, we built a corpus of 3 000 figures that allows us to reliably estimate the numbers and types of images in biomedical articles. These figures were drawn randomly from the open access subset of PubMed Central and then manually annotated. They were split into subfigures when the figure consisted of several components. Figure 1 shows the resulting categories and subcategories. This classification scheme is based on five basic image categories: Experimental/Microscopy, Graph, Diagram, Clinical and Picture, each divided into multiple subcategories. It shows that bar graphs (12.4%), black-on-white gels (12.0%), fluorescence microscopy images (9.4%), and line graphs (8.1%) are the most frequent subfigure types (all percentages are relative to the entire set of images).

We targeted different kinds of graphs (i.e., diagrams with axes) in previous work [7], and we decided to focus this work on the second most common

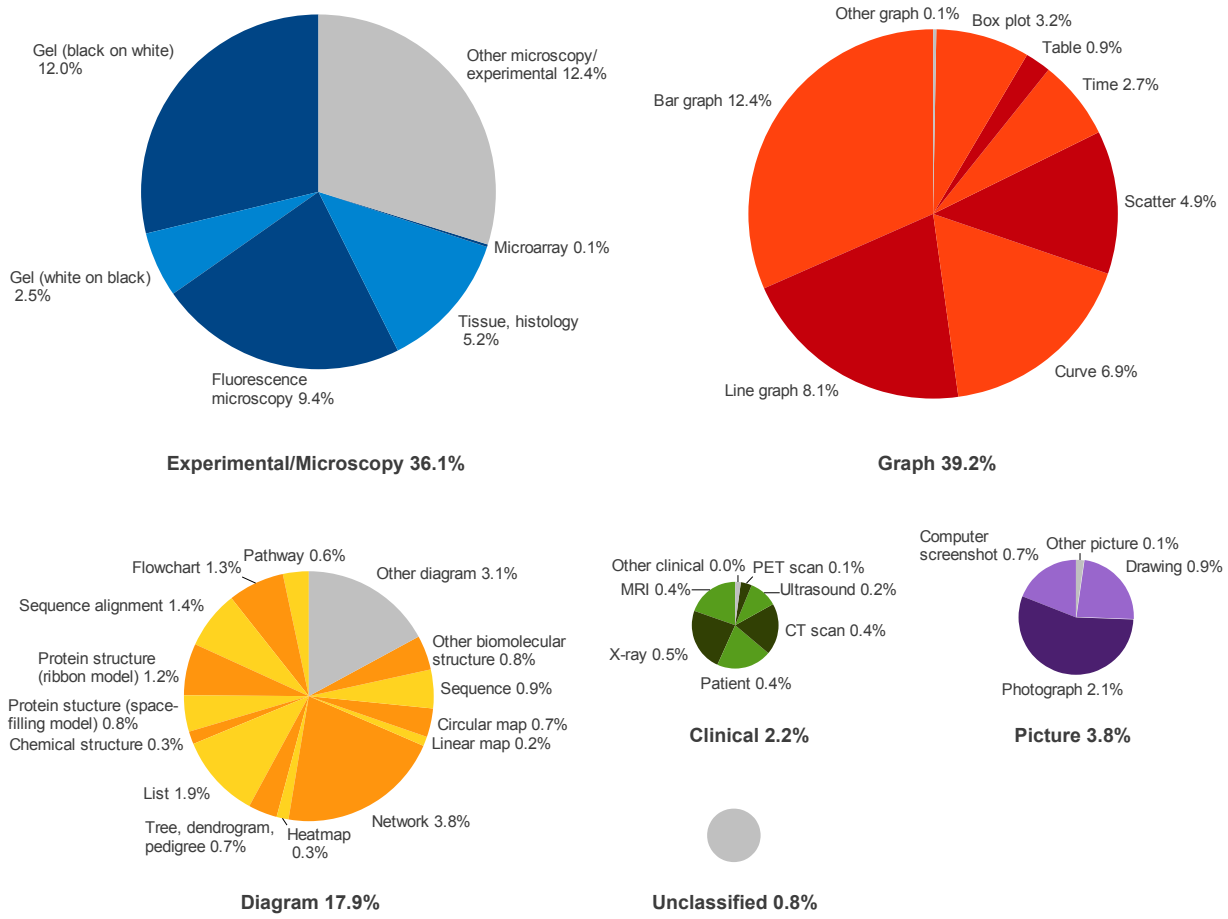

Figure 1: Categorization of images from open access articles of PubMed Central.

type of images: gel diagrams. They are the result of gel electrophoresis, which is a common method to analyze DNA, RNA and proteins. Southern, Western and Northern blotting [8–10] are among the most common applications of gel electrophoresis. The resulting experimental artifacts are often shown in biomedical publications in the form of gel images as evidence for the discussed findings such as protein-protein interactions or protein expressions under different conditions. Often, not all details of the results shown in these images are explicitly stated in the caption or the article text. For these reasons, it would be of high value to be able to reliably mine the relations encoded in these images.

A closer look at gel images reveals that they follow regular patterns to encode their semantic relations. Figure 2 shows two typical examples of gel images together with a table representation of the involved relations. The ultimate objective of our approach (for which we can only present a partial solution here) is to automatically extract at least some of these relations from the respective images, possibly in conjunction with classical text mining techniques. The first example shows a Western blot for detecting two proteins (14-3-3σ and β-actin as a control) in four different cell lines (MDA-MB-231, NHEM, C8161.9, and LOX, the first of which is used as a control). There are two rectangular gel segments arranged in a way to form a 2×4 grid for the individual eight measurements combining each protein with each cell line. A gel diagram can be considered a kind of matrix with pictures of experimental ar-

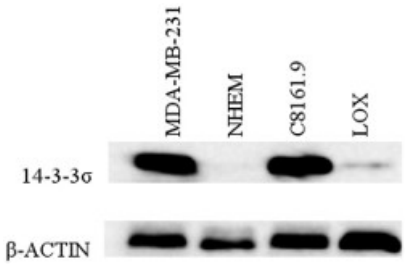

Condition	Measurement	Result
MDA-MB-231	14-3-3σ	high expression
NHEM	14-3-3σ	no expression
C8161.9	14-3-3σ	high expression
LOX	14-3-3σ	low expression
MDA-MB-231	β-actin	high expression
NHEM	β-actin	high expression
C8161.9	β-actin	high expression
LOX	β-actin	high expression

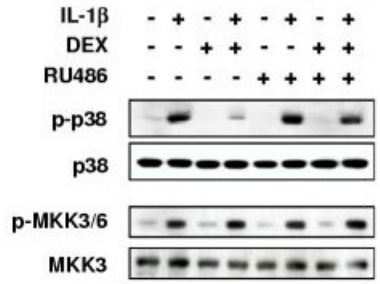

Condition			Measurement	Result
IL-1β (–)	DEX (–)	RU486 (–)	p-p38	low expression
IL-1β (+)	DEX (–)	RU486 (–)	p-p38	high expression
IL-1β (–)	DEX (+)	RU486 (–)	p-p38	no expression
IL-1β (+)	DEX (+)	RU486 (–)	p-p38	low expression
IL-1β (–)	DEX (–)	RU486 (+)	p-p38	no expression
IL-1β (+)	DEX (–)	RU486 (+)	p-p38	high expression
IL-1β (–)	DEX (+)	RU486 (+)	p-p38	low expression
IL-1β (+)	DEX (+)	RU486 (+)	p-p38	high expression
...			...	...

Figure 2: Two examples of gel images from biomedical publications (PMID 19473536 and 15125785) with tables showing the relations that could be extracted from them

tifacts as content. The tables to the right illustrate the semantic relations encoded in the gel diagrams. Each relation instance consists of a condition, a measurement and a result. The proteins are the entities being measured under the conditions of the different cell lines. The result is a certain degree of expression indicated by the darkness of the spots (or brightness in the case of white-on-black gels). The second example is a slightly more complex one. Several proteins are tested against each other in a way that involves more than two dimensions. In this case, the use of "+" and "–" labels is a frequent technique to denote the different possible combinations of a number of conditions. Apart from that, the principles are the same. In this case, however, the number of relations is much larger. Only the first eight of a total of 32 relation instances are shown in the table to the right. In such cases, the text rarely mentions all these relations in an explicit way, and the image is therefore the only accessible source.

Background

In principle, image mining involves the same processes as classical literature mining [11]: document categorization, named entity tagging, fact extraction, and collection-wide analysis. However, there are some subtle differences. Document categorization corresponds to image categorization, which is different in the sense that it has to deal with features based on the two-dimensional space of pixels, but otherwise the same principles of automatic categorization apply. Named entity tagging is different in two ways: pinpointing the mention of an entity is more difficult with images (a large number of pixels versus a couple of characters), and OCR errors have to be considered. Fact extraction in classical literature mining involves the analysis of the syntactic structure of the sentences. In images, in contrast, there are rarely complete sentences, but the semantics is rather encoded by graphical means. Thus, instead of parsing sentences, one has to analyze graphical elements and their relation to each other. The last process, collection-wide analysis, is a higher-level problem, and therefore no fundamental differences can be expected. Thus, image mining builds upon the same general stages as classical text mining, but with some subtle yet important differences.

Image mining on biomedical publications is not a new idea. It has been applied for the extraction of subcellular location information [12], the detec-

tion of panels of fluorescence microscopy images [13], the extraction of pathway information from diagrams [14], and the detection of axis diagrams [7]. Also, there is a large amount of existing work on how to process gel images [15–19] and databases have been proposed to store the results of gel analyses [20]. These techniques, however, take as input plain gel images, which are not readily accessible from biomedical papers, because they make up just parts of the figures. Furthermore, these tools are designed for researchers who want to analyze their gel images and not to read gel diagrams that have already been analyzed and annotated by a researcher. Therefore, these approaches do not tackle the problem of recognizing and analyzing the labels of gel images. Some attempts to classify biomedical images include gel figures [21], which is, however, just the first step in locating them and analyzing their labels and their structure. To our knowledge, nobody has yet tried to perform image mining on gel diagrams.

Approach and Methods

Figure 3 shows the procedure of our approach to image mining from gel diagrams. It consists of seven steps: figure extraction, segmentation, text recognition, gel detection, gel panel detection, named entity recognition and relation extraction.[1]

Using structured article representations, the first step is trivial. For steps two and three, we rely on existing work. The main focus of this paper lies on steps four and five: the detection of gels and gel panels. In the discussion section, we present some preliminary results on step six of recognizing named entities, and sketch how step seven could be implemented, for which we cannot provide a concrete solution at this point.

To practically evaluate our approach, we ran our pipeline on the entire open access subset of PubMed Central (though not all figures made it through the whole pipeline due to technical difficulties).

Figure Extraction

A large portion of the articles of the open access subset of the PubMed Central database are available as structured XML files with additional image files for the figures. We only use these articles so far, which makes the figure extraction task very easy. It would be more difficult, though definitely feasible, to extract the figures from PDF files or even bitmaps of scanned articles (see [22] and http://pdfjailbreak.com for approaches on extracting the structure of articles in PDF format).

Segmentation and Text Recognition

For the next two steps — segment detection and subsequent text recognition —, we rely on our previous work [23,24]. This method includes the detection of layout elements, edge detection, and text recognition with a novel pivoting approach. For optical character recognition (OCR), the Microsoft Document Imaging package is used, which is available as part of Microsoft Office 2003. Overall, this approach has been shown to perform better than other existing approaches for the images found in biomedical publications [23]. We do not go into the details here, as this paper focuses on the subsequent steps.

Due to some limitations of the segmentation algorithm when it comes to rectangles with low internal contrast (like gels), we applied a complementary very simple rectangle detection algorithm.

Gel Segment Detection

Based on the results of the above-mentioned steps, we try to identify gel segments. Such gel segments typically have rectangular shapes with darker spots on a light gray background, or — less commonly — white spots on a dark background. We decided to use machine learning techniques to generate classifiers to detect such gel segments. To do so, we defined 39 numerical features for image segments: the coordinates of the relative position (within the image), the relative and absolute width and height, 16 grayscale histogram features, three color features (for red, green and blue), 13 texture features (coarseness, presence of ripples, etc.) based on [25], and the number of recognized characters.

To train the classifiers, we took a random sample of 500 figures, for which we manually annotated the gel segments. In the same way, we obtained a second sample of another 500 figures for testing the

[1] Due to the fact that many figures consist of multiple panels of different types, we go straight to gel segment detection without first classifying entire images. Most gel panels share their figure with other panels, which makes automated classification difficult at the image level.

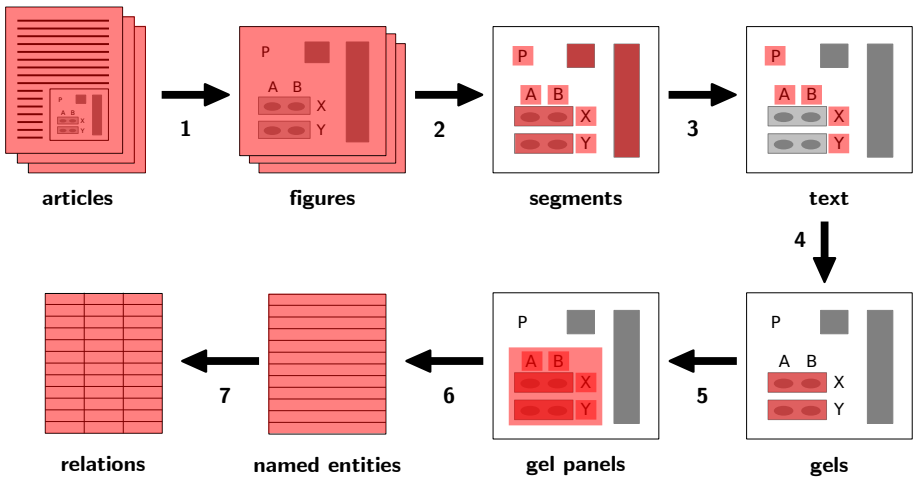

Figure 3: The procedure of our approach: (1) figure extraction, (2) segmentation, (3) text recognition, (4) gel detection, (5) gel panel detection, (6) named entity recognition, and (7) relation extraction.

classifiers.[2] We used the Weka toolkit and opted for random forest classifiers based on 75 random trees. Using different thresholds to adjust the trade-off between precision and recall, we generated a classifier with good precision and another one with good recall. Both of them are used in the next step. We tried other types of classifiers including naive Bayes, Bayesian networks [26], PART decision lists [27], and convolutional networks [28], but we achieved the best results with random forests.

Gel Panel Detection

A gel panel typically consists of several gel segments and comes with labels describing the involved genes, proteins, and conditions. For our goal, it is not sufficient to just detect the figures that contain gel panels, but we also have to extract their positions within the figures and to access their labels. This is not a simple classification task, and therefore machine learning techniques do not apply that easily. For that reason, we used a detection procedure based on hand-coded rules.

In a first step, we group gel segments to find contiguous gel regions that form the center part of gel panels. To do so, we start with looking for segments that our high-precision classifier detects as gel segments. Then, we repeatedly look for adjacent gel segments, this time applying the high-recall classifier, and merge them. Two segments are considered neighbors if they are at most 50 pixels apart[3] and do not have any text segment between them. Thus, segments which could be gel segments according to the high-recall classifier make it into a gel panel only if there is at least one high-precision segment in their group. The goal is to detect panels with high precision, but also to detect the complete panels and not just parts of them. We focus here on precision because low recall can be leveraged by the large number of available gel images. Furthermore, as the open access part of PubMed Central only makes up a small subset of all biomedical publications, recall in a more general sense is anyway limited by the proportion of open access publications.

As a next step, we collect the labels in the form of text segments located around the detected gel regions. For a text segment to be attributed to a certain gel panel, its nearest edge must be at most 30 pixels away from the border of the gel region and its farthest edge must not be more than 150 pixels

[2]We double-checked these manual annotations to check their quality, which revealed only four misclassified segments in total for the training and test samples (0.016% of all segments).

[3]We are using absolute distance values at this point. A more refined algorithm could apply some sort of relative measure. However, the resolution of the images does not vary that much, which is why absolute values worked out well so far.

away. We end up with a representation of a gel panel consisting of two parts: a center region containing a number of gel segments and a set of labels in the form of text segments located around the center region.

To evaluate this algorithm, we collected yet another sample of 500 figures, in which 106 gel panels in 61 different figures were revealed by manual annotation.[4] Based on this sample, we manually checked whether our algorithm is able to detect the presence and the (approximate) position of the gel panels.

Results

The top part of Table 1 shows the result of the gel detection classifier. We generated three different classifiers from the training data, one for each of the threshold values 0.15, 0.3 and 0.6. Lower threshold values lead to higher recall at the cost of precision, and vice versa. In the balanced case, we achieved an F-score of 75%. To get classifiers with precision or recall over 90%, F-score goes down significantly, but stays in a sensible range. These two classifiers (thresholds 0.15 and 0.6) are used in the next step. To interpret these values, one has to consider that gel segments are greatly outnumbered by non-gel segments. Concretely, only about 3% are gel segments. More sophisticated accuracy measures for classifier performance, such as the area under the ROC curve [29], take this into account. For the presented classifiers, the area under the ROC curve is 98.0% (on a scale from 50% for a trivial, worthless classifier to 100% for a perfect one).

The results of the gel panel detection algorithm are shown in the bottom part of Table 1. The precision is 95% at a recall of 37%, leading to an F-score of 53%. The comparatively low recall is mainly due to the general problem of pipeline-based approaches that the various errors from the earlier steps accumulate and are hard to correct at a later stage in the pipeline.

Table 2 shows the results of running the pipeline on PubMed Central. We started with about 410 000 articles, the entire open access subset of PubMed Central at the time we downloaded them (February 2012). We successfully parsed the XML files of 94% of these articles (for the remaining articles, the XML file was missing or not well-formed, or other unexpected errors occurred). The successful articles contained around 1 100 000 figures, for some of which our segment detection step encountered image formatting errors or other internal errors, or was just not able to detect any segments. We ended up with more than 880 000 figures, in which we detected about 86 000 gel panels, i.e. roughly ten out of 100 figures. For each of them, we found on average 3.6 labels with recognized text. After tokenization, we identified about 76 000 gene names in these gel labels, which corresponds to 6.8% of the tokens. Considering all text segments (including but not restricted to gel labels), only 3.3% of the tokens are detected as gene names.[5]

Discussion

The presented results show that we are able to detect gel segments with high accuracy, which allows us to subsequently detect whole gel panels at a high precision. The recall of the panel detection step is relatively low, but with about 37% still in a reasonable range. As mentioned above, we can leverage the high number of available figures, which makes precision more important than recall. Running our pipeline on the whole set of open access articles from PubMed Central, we were able to retrieve 85 942 potential gel panels (around 95% of which we can expect to be correctly detected).

The next step would be to recognize the named entities mentioned in the gel labels. To this aim, we did a preliminary study to investigate whether we are able to extract the names of genes and proteins from gel diagrams. To do so, we tokenized the label texts and looked for entries in the Entrez Gene database to match the tokens. This look-up was done in a case-sensitive way, because many names in gel labels are acronyms, where the specific capitalization pattern can be critical to identify the respective entity. We excluded tokens that have less than three characters, are numbers (Arabic or Latin), or correspond to common short words (retrieved from a list of the 100 most frequent words in biomedical articles). In addition, we extended this exclusion list with 22 general words that are frequently used in the context of gel diagrams, some of which co-

[4] Again, these manual annotations were double-checked to ensure their quality. Five errors were found and fixed in this process.

[5] The low numbers are partially due to the fact that a considerable part of the tokens are "junk tokens" produced by the OCR step when trying to recognize characters in segments that do not contain text.

	Method	Threshold	Precision	Recall	F-score	ROC area
Segments	Random forests	0.15	0.439	0.909	0.592	0.980
		0.30	0.765	0.739	0.752	
		0.60	0.926	0.301	0.455	
	Naive Bayes		0.172	0.739	0.279	0.883
	Bayesian network		0.394	0.531	0.452	0.914
	PART decision list		0.631	0.496	0.555	0.777
	Convolutional networks		0.142	0.949	0.248	
Panels	Hand-coded rules		0.951	0.368	0.530	

Table 1: The results of the gel segment detection classifiers (top) and the gel panel detection algorithm (bottom)

Total articles	410 950
Processed articles	386 428
Total figures from processed articles	1 110 643
Processed figures	884 152
Detected gel panels	85 942
Detected gel panels per figure	0.097
Detected gel labels	309 340
Detected gel labels per panel	3.599
Detected gene tokens	1 854 609
Detected gene tokens in gel labels	75 610
Gene token ratio	0.033
Gene token ratio in gel labels	0.068

Table 2: The results of running the pipeline on the open access subset of PubMed Central

incide with gene names according to Entrez.[6] Since gel electrophoresis is a method to analyze genes and proteins, we would expect to find more such mentions in gel labels than in other text segments of a figure. By measuring this, we get an idea of whether the approach works out or not. In addition, we manually checked the gene and protein names extracted from gel labels after running our pipeline on 2 000 random figures. In 124 of these figures, at least one gel panel was detected. Table 3 shows the results of this preliminary evaluation. Almost two-thirds of the detected gene/protein tokens (65.3%) were correctly identified. 9% thereof were correct but could be more specific, e.g. when only "actin" was recognized for "β-actin" (which is not incorrect but of course much harder to map to a meaningful identifier). The incorrect cases (34.6%) can be split into two classes of roughly the same size: some recognized tokens were actually not mentioned in the figure but emerged from OCR errors; other tokens were correctly recognized but incorrectly classified as gene or protein references. Although there is certainly much room for improvement, this simple gene detection step seems to perform reasonably well.

For the last step, relation extraction, we cannot present any concrete results at this point. After recognizing the named entities, we would have to disambiguate them, identify their semantic roles (condition, measurement or something else), align the gel images with the labels, and ultimately quantify the degree of expression. To improve the quality of the results, combinations with classical text mining techniques should be considered. This is all future work. We expect to be able to profit to a large extent from existing work to disambiguate protein and gene names [30, 31] and to detect and analyze gel spots (see the existing work mentioned above).

It seems reasonable to assume that these results

[6]These words are: *min*, *hrs*, *line*, *type*, *protein*, *DNA*, *RNA*, *mRNA*, *membrane*, *gel*, *fold*, *fragment*, *antigen*, *enzyme*, *kinase*, *cleavage*, *factor*, *blot*, *pro*, *pre*, *peptide*, and *cell*.

	absolute	relative
Total	**156**	**100.0%**
Incorrect	**54**	**34.6%**
– Not mentioned (OCR errors)	28	17.9%
– Not references to genes or proteins	26	16.7%
Correct	**102**	**65.3%**
– Partially correct (could be more specific)	14	9.0%
– Fully correct	88	56.4%

Table 3: Numbers of recognized gene/protein tokens in 2 000 random figures

can be combined with existing techniques of term disambiguation and gel spot detection at a satisfactory level of accuracy. We plan to investigate this in future work.

As mentioned above, we have started to investigate how the gel segment detection step could be improved by the use of the image recognition technique of convolutional networks (ConvNet) [28]. We started with a simplified approach to the one presented in [32]. In this approach, images are tiled into small quadratic pieces. We used a single network (and not several parallel networks), based on 48×48 input tile images with three layers of convolutions. The first layer takes eight 5×5 convolutions and is followed by a 2×2 sub-sampling. The second layer takes twenty four 5×5 convolutions and is followed by a 3×3 sub-sampling. The last layer takes seventy two 6×6 convolutions, which leads to a fully connected layer. We trained our ConvNet on the 500 images of the training set, where we manually annotated the tiles as *gel* and *non-gel*. With the use of EBLearn [33], this trained ConvNet classified the tiles of the 500 images of our testing set. The classified tiles can then be reconstructed into a mask image, as shown in Figure 4. A manual check of the clusters of recognized gel tiles led to the results shown in Table 1. Recall is very good (95%) but precision is very poor (14%), leading to an F-score of 25%. This is much worse than the results we got with our random forest approach, which is why ConvNet is currently not part of our pipeline. We hope, however, that we can further optimize this ConvNet approach and combine it with random forests to exploit their (hopefully) complementary benefits. Using ConvNet to classify complete images as *gel-image* or *non-gel-image* and adjusting the classification to account for unbalanced classes, we were able to obtain an F-score of 74%, which makes us confident that a combination of the two approaches could lead to a significant improvement of our gel segment detection step. As an alternative approach, we will try to run ConvNet on down-scaled entire panels rather than small tiles, as described in [34]. Furthermore, we will experiment with parallel networks instead of single ones to improve accuracy.

The results obtained from our gel recognition pipeline indicate that it is feasible to extract relations from gel images, but it is clear that this procedure is far from perfect. The automatic analysis of bitmap images seems to be the only efficient way to extract such relations from existing publications, but other publishing techniques should be considered for the future. The use of vector graphics instead of bitmaps would already greatly improve any subsequent attempts of automatic analysis. A further improvement would be to establish accepted standards for different types of biomedical diagrams in the spirit of the Unified Modeling Language, a graphical language widely applied in software engineering since the 1990s. Ideally, the resulting images could directly include semantic relations in a formal notation, which would make relation mining a trivial procedure. If authors are supported by good tools to draw diagrams like gel images, this approach could turn out to be feasible even in the near future.

Concretely, we would like to take the opportunity to postulate the following actions, which we think should be carried out to make the content of images in biomedical articles more accessible:

- **Stop pressing diagrams into bitmaps!** Unless the image only consists of one single photograph, screenshot, or another kind of picture that only has bitmap representation, vector graphics should be used for article figures.

- **Let data and metadata travel from the**

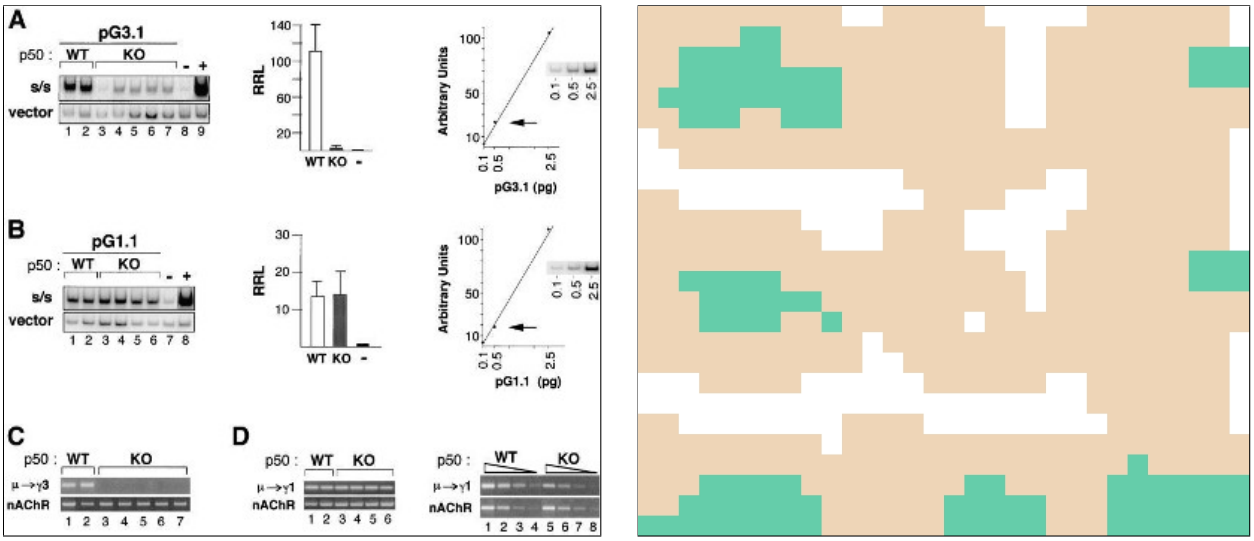

Figure 4: Original and mask image after ConvNet classification for an exemplary image from PMID 14993249. Green means *gel*; brown means *other*; and white means *not enough gradient information*.

tools that generate diagrams to the final articles! Whenever the specific tool that is used to generate the diagram "knows" that a certain graphical element refers to an organism, a gene, an interaction, a point in time, or another kind of entity, then this information should be stored in the image file, passed on, and finally published with the article.

- **Use RDF vocabularies to embed semantic annotations in diagrams!** Tools for creating scientific diagrams should use RDF notation and stick to existing standardized schemas (or define new ones if required) to annotate the diagram files they create.
- **Define standards for scientific diagrams!** In the spirit of the Unified Modeling Language, the biomedical community should come up with standards that define the appearance and meaning of different types of diagrams.

Obviously, different groups of people need to be involved in these actions, namely article authors, journal editors, and tool developers. It is relatively inexpensive to follow these postulates (though it might require some time), which in turn would greatly improve data sharing, image mining, and scientific communication in general. Standardized diagrams could be the long sought solution to the problem of how to let authors publish computer-processable formal representations for (part of) their results. This can build upon the efforts of establishing an open annotation model [35, 36].

Conclusions

Successful image mining from gel diagrams in biomedical publications would unlock a large amount of valuable data. Our results show that gel panels and their labels can be detected with high accuracy, applying machine learning techniques and hand-coded rules. We also showed that genes and proteins can be detected in the gel labels with satisfactory precision.

Based on these results, we believe that this kind of image mining is a promising and viable approach to provide more powerful query interfaces for researchers, to gather relations such as protein-protein interactions, and to generally complement existing text mining approaches. At the same time, we believe that an effort towards standardization of scientific diagrams such as gel images would greatly improve the efficiency and precision of image mining at relatively low additional costs at the time of publication.

Competing Interests

The authors declare that they have no competing interests.

Authors' Contributions

TK was the main author and main contributor of the presented work. He was responsible for designing and implementing the pipeline, gathering the data, performing the evaluation, and analyzing the results. MLN applied, trained, and evaluated the ConvNet classifier, and contributed to the annotation of the test sets. TL built and analyzed the corpus for the preparatory evaluation. MK contributed to the conception and the design of the approach and to the analysis of the results. All authors have been involved in drafting or revising the manuscript, and all authors read and approved the final manuscript.

Acknowledgments

This study has been supported by the National Library of Medicine grant 5R01LM009956.

References

1. Yu H, Lee M: **Accessing bioscience images from abstract sentences**. *Bioinformatics* 2006, **22**(14):e547–e556.
2. Zweigenbaum P, Demner-Fushman D, Yu H, Cohen KB: **Frontiers of biomedical text mining: current progress**. *Briefings in Bioinformatics* 2007, **8**(5):358–375.
3. Peng H: **Bioimage informatics: a new area of engineering biology**. *Bioinformatics* 2008, **24**(17):1827–1836, [http://dx.doi.org/10.1093/bioinformatics/btn346].
4. Xu S, McCusker J, Krauthammer M: **Yale Image Finder (YIF): a new search engine for retrieving biomedical images**. *Bioinformatics* 2008, **24**(17):1968–1970, [http://dx.doi.org/10.1093/bioinformatics/btn340].
5. Hearst MA, Divoli A, Guturu H, Ksikes A, Nakov P, Wooldridge MA, Ye J: **BioText Search Engine**. *Bioinformatics* 2007, **23**(16):2196–2197, [http://dx.doi.org/10.1093/bioinformatics/btm301].
6. Kuhn T, Krauthammer M: **Image Mining from Gel Diagrams in Biomedical Publications**. In *Proceedings of the 5th International Symposium on Semantic Mining in Biomedicine (SMBM 2012)*, University of Zurich (Zurich, Switzerland) 2012:26–33, [http://www.zora.uzh.ch/64476/].
7. Kuhn T, Luong T, Krauthammer M: **Finding and Accessing Diagrams in Biomedical Publications**. In *Proceedings of the American Medical Informatics Association (AMIA) 2012 Annual Symposium*, American Medical Informatics Association (Bethesda, MD, USA) 2012.
8. Southern E: **Detection of specific sequences among DNA fragments separated by gel electrophoresis**. *J. Mol. Biol.* 1975, **98**(3):503–517.
9. Alwine JC, Kemp DJ, Stark GR: **Method for detection of specific RNAs in agarose gels by transfer to diazobenzyloxymethyl-paper and hybridization with DNA probes**. *Proceedings of the National Academy of Sciences* 1977, **74**(12):5350.
10. Burnette WN: **Western Blotting: Electrophoretic transfer of proteins from sodium dodecly sulfate-polyacrylamide gels to unmodified nitrocellulose and radiographic detection with antibody and radioiodinated protein A**. *Anal. Biochem.* 1981, **112**:195–203.
11. De Bruijn B, Martin J: **Getting to the (c)ore of knowledge: mining biomedical literature**. *International Journal of Medical Informatics* 2002, **67**(1-3):7–18.
12. Murphy RF, Kou Z, Hua J, Joffe M, Cohen WW: **Extracting and Structuring Subcellular Location Information from On-line Journal Articles: The Subcellular Location Image Finder**. In *Proceedings of the IASTED International Conference on Knowledge Sharing and Collaborative Engineering (KSCE-2004)*, ACTA Press (Calgary, AB, Canada) 2004:109–114, [http://www.actapress.com/Abstract.aspx?paperId=17244].
13. Qian Y, Murphy RF: **Improved recognition of figures containing fluorescence microscope images in online journal articles using graphical models**. *Bioinformatics* 2008, **24**(4):569–576, [http://dx.doi.org/10.1093/bioinformatics/btm561].
14. Kozhenkov S, Baitaluk M: **Mining and integration of pathway diagrams from imaging data**. *Bioinformatics* 2012, **28**(5):739–742, [http://dx.doi.org/10.1093/bioinformatics/bts018].
15. Lemkin PF: **Comparing two-dimensional electrophoretic gel images across the Internet**. *Electrophoresis* 1997, **18**(3-4):461–470.
16. Luhn S, Berth M, Hecker M, Bernhardt J: **Using standard positions and image fusion to create proteome maps from collections of two-dimensional gel electrophoresis images**. *Proteomics* 2003, **3**(7):1117–1127.
17. Cutler P, Heald G, White IR, Ruan J: **A novel approach to spot detection for two-dimensional gel electrophoresis images using pixel value collection**. *Proteomics* 2003, **3**(4):392–401.
18. Rogers M, Graham J, Tonge RP: **Statistical models of shape for the analysis of protein spots in two-dimensional electrophoresis gel images**. *Proteomics* 2003, **3**(6):887–896.
19. Zerr T, Henikoff S: **Automated band mapping in electrophoretic gel images using background information**. *Nucleic Acids Research* 2005, **33**(9):2806–2812.
20. Schlamp K, Weinmann A, Krupp M, Maass T, Galle P, Teufel A: **BlotBase: a northern blot database**. *Gene* 2008, **427**(1-2):47–50.
21. Rodriguez-Esteban R, Iossifov I: **Figure mining for biomedical research**. *Bioinformatics* 2009, **25**(16):2082–2084, [http://dx.doi.org/10.1093/bioinformatics/btp318].
22. Ramakrishnan C, Patnia A, Hovy EH, Burns GA, et al.: **Layout-aware text extraction from full-text PDF of scientific articles**. *Source code for biology and medicine* 2012, **7**:255–258.
23. Xu S, Krauthammer M: **A new pivoting and iterative text detection algorithm for biomedical images**. *J. of Biomedical Informatics* 2010, **43**(6):924–931, [http://dx.doi.org/10.1016/j.jbi.2010.09.006].
24. Xu S, Krauthammer M: **Boosting text extraction from biomedical images using text region detection**. In *Biomedical Sciences and Engineering Conference (BSEC), 2011*, IEEE (New York City, NY, USA) 2011:1–4.
25. Haralick RM, Shanmugam K, Dinstein I: **Textural Features for Image Classification**. *IEEE Transactions on Systems, Man, and Cybernetics* 1973, **3**(6):610–621, [http://dx.doi.org/10.1109/TSMC.1973.4309314].
26. Cooper GF, Herskovits E: **A Bayesian method for the induction of probabilistic networks from data**. *Machine learning* 1992, **9**(4):309–347.
27. Frank E, Witten IH: **Generating Accurate Rule Sets Without Global Optimization**. In *Proceedings of the Fifteenth International Conference on Machine Learning*, Morgan Kaufmann Publishers (Burlington, MA, USA) 1998:144–151.

28. LeCun Y, Bengio Y: **Convolutional networks for images, speech, and time series**. *The handbook of brain theory and neural networks* 1995, **3361**.

29. Bradley AP: **The use of the area under the ROC curve in the evaluation of machine learning algorithms**. *Pattern recognition* 1997, **30**(7):1145–1159.

30. Tanabe L, Wilbur WJ: **Tagging gene and protein names in biomedical text**. *Bioinformatics* 2002, **18**(8):1124–1132.

31. Lu Z, Kao HY, Wei CH, Huang M, Liu J, Kuo CJ, Hsu CN, Tsai R, Dai HJ, Okazaki N, et al.: **The gene normalization task in BioCreative III**. *BMC bioinformatics* 2011, **12**(Suppl 8):S2.

32. Barbano PE, Nagy ML, Krauthammer M: **Energy-Based Architecture for Classification of Publication Figures**. In *Proceedings of the Biomedical Science and Engineering Center Conference (BSEC 2013)*, IEEE (New York City, NY, USA) 2013.

33. Sermanet P, Kavukcuoglu K, LeCun Y: **EBlearn: Open-source energy-based learning in C++**. In *Proceedings of the 21st International Conference on Tools with Artificial Intelligence (ICTAI'09)*, IEEE (New York City, NY, USA) 2009:693–697.

34. Krizhevsky A, Sutskever I, Hinton G: **ImageNet classification with deep convolutional neural networks**. *Advances in Neural Information Processing Systems* 2012, **25**:1106–1114.

35. Ciccarese P, Ocana M, Garcia Castro LJ, Das S, Clark T: **An open annotation ontology for science on web 3.0**. *J Biomed Semantics* 2011, **2**(Suppl 2):S4.

36. Sanderson R, Ciccarese P, Van de Sompel H: **Open Annotation Data Model**. Community draft, W3C 2013, [http://www.openannotation.org/spec/core/20130208/index.html].